

Multi Objective Dyeing Optimisation

Key Performance Indicators

—14.2%

water drawn per batch

—11.6%

thermal energy per batch

—13.8%

salt and electrolyte dosed

+8.7%

right first time shade pass rate

—12.9%

process carbon per kilogram
died



25 runs

a compact design of experiments is enough to map the process response across dye, salt, alkali, temperature and time before optimisation begins



0.96

surrogate accuracy, as R^2 , predicting exhaustion, fixation and colour strength from the recipe inputs, so recipes are tested in software first




under 14 s

for the agent to converge on a balanced recipe across every competing objective, faster than the population based methods it replaces

Reactive dyeing of cotton is one of the most resource intensive stages in textile production. Achieving a target shade depends on the interaction of dye dosage, salt concentration, alkali, liquor ratio, temperature and time, and small deviations push a batch off shade and force costly redyeing or shade correction. Each redye repeats the water, salt, energy and effluent load of the original batch. Operators must hold shade accuracy while also containing water draw, thermal energy, electrolyte use and cycle time, and these aims pull against one another. Conventional recipe setting leans on standing recipes and colourist judgement, which leaves a wide gap between the best and the typical batch and offers no systematic way to trade one objective against another.

EntroMetrix builds a multi objective optimiser on the operator's own historical batch records, laboratory dips and continuous plant logs. Random forest and Gaussian process surrogates learn how recipe and schedule inputs map to exhaustion, fixation, colour strength and the resource cost of each batch. A reinforcement learning agent then searches this learned response, proposing recipes and dosing schedules that hold the target shade while lowering water, salt, chemical and energy use. Rather than returning a single answer, the agent learns a policy across the competing objectives and surfaces a Pareto front, so the plant can select an operating point that suits its current priorities on cost, throughput and effluent.

Inside the optimiser



The optimiser weighs shade accuracy against the resource and energy cost of a batch, on the left, and learns a policy that balances them, on the right. The chosen operating point sits at the knee of the Pareto front, where a small relaxation in cost buys most of the available shade accuracy.

What the Optimiser Does

- **Learns from plant data:** trains on existing batch records, dips and plant logs, with no new trials needed to start.
- **Surrogate models:** random forest and Gaussian process models predict exhaustion, fixation, colour strength and resource cost.
- **Reinforcement learning:** an agent searches the learned response to balance shade accuracy against water, salt and energy.
- **Pareto front:** returns a range of balanced recipes rather than one, so the plant chooses its own operating point.
- **Recipe and schedule:** proposes dye, salt and alkali dosing with temperature and time profiles for each shade.
- **Right first time:** targets shade within tolerance on the first pass to cut redyes and shade correction.




Figure 1. Water, energy, salt and process carbon per kilogram dyed, baseline against optimised. One policy moves every objective in the same direction rather than trading one away for another.

The optimiser is built in weeks because it begins with data the plant already holds. Historical batch cards, dyebath dosing logs, laboratory dip readings and colour measurements are enough to characterise the process, together with the standing recipes and the operating limits captured with the dyehouse team. Key assumptions are documented and confirmed by the plant's colourists before any recommendation is made.

EntroMetrix delivered this optimiser on a reactive dyeing and finishing operation. Surrogate models learned exhaustion, fixation, colour strength and shade from the recipe inputs to an accuracy of 0.96 on held back batches, and a compact design of experiments of twenty five runs was enough to map the process response before optimisation began. Against that response the agent converged on a balanced recipe in under fourteen seconds, faster than the population based multi objective methods it replaces.

Validation is retrospective and blind. Recommendations are checked against past batches, using success criteria agreed in advance, and every candidate recipe is confirmed on a laboratory dip before it reaches bulk production. Optimisation stays inside the plant's proven process window and dosing constraints, and the colourist reviews and approves each recipe, so judgement stays in the process. As new batches are dyed they feed back to sharpen the surrogates and refine the policy.

What the Optimiser Does (cont.)

- **Retrospective test:** recommendations are checked against past batches before any recipe reaches the floor.
- **Safe operating limits:** optimisation stays within the plant's proven process window and dosing constraints.
- **Laboratory dip check:** candidate recipes are confirmed on lab dips ahead of bulk production.
- **Colourist in the loop:** colourists review and approve proposed recipes, keeping judgement in the process.
- **Continuous learning:** new batches feed back to sharpen the surrogates and refine the policy.
- **Effluent tracking:** salt, chemical and water reductions are logged against baseline for reporting.

What the Optimiser Delivers




Figure 2. First pass shade pass rate against redye share, baseline and optimised. Holding shade on the first pass is what turns recipe quality into avoided water, salt and energy.




Figure 3. Batch cycle time by shade depth. Tighter recipes and schedules take time out of every class, freeing capacity from the same plant.




Figure 4. Shade error across a run of batches. The optimised policy holds colour inside tolerance batch after batch, closing the gap between the best and the typical batch.

Shade Right First Time, at Lower Cost

Delivered on a reactive dyeing and finishing operation

Mohammed Ali, CEO · mo@entrometrix.ai · +44 (0) 773 800 0590